

# **Exhibit 23 to Amended Complaint**

## **Intellectual Ventures I LLC and Intellectual Ventures II LLC**

**Example American Count IX Systems and Services  
U.S. Patent No. 8,352,584 (“the ’584 Patent”)**

The Accused Systems and Services include without limitation American systems and services that utilize Kubernetes; all past, current, and future systems and services that operate in the same or substantially similar manner as the specifically identified systems and services; and all past, current, and future American systems and services that have the same or substantially similar features as the specifically identified systems and services (“Example American Count IX Systems and Services” or “American Systems and Services”).<sup>1</sup>

On information and belief, the American Systems and Services use Kubernetes in public and/or private cloud(s). For example, American posts, or has posted, job opportunities that require familiarity with Kubernetes containerization concepts.

- Example of job listing for an Engineer/Sr Engineer, IT Situational Awareness at American Airlines which requires proficiency in Kubernetes. [https://jobs.aa.com/job/EngineerSr-Engineer%2C-IT-Situational-Awareness/75837-en\\_US](https://jobs.aa.com/job/EngineerSr-Engineer%2C-IT-Situational-Awareness/75837-en_US).
- Example of job listing for an Associate Developer, IT Applications at American Airlines which requires proficiency in Kubernetes. [https://jobs.aa.com/job/Associate-Developer%2C-IT-Applications/75816-en\\_US](https://jobs.aa.com/job/Associate-Developer%2C-IT-Applications/75816-en_US).
- Example of Senior Java Full Stack Developer position at American Airlines which mentions use of Kubernetes. <https://www.linkedin.com/in/rohitha-m6363/>.
- Example of Sr. Kubernetes Engineer/Architect position at American Airlines which mentions heavy use of Kubernetes. <https://www.linkedin.com/in/sudheer-patchari/>.
- Example of Kubernetes Engineer position at American Airlines which mentions heavy use of Kubernetes. <https://www.linkedin.com/in/sridhar-pulluri-199b56250/>.
- Example of Sr. Cloud Infra DevSecOps Engineer/Architect position at American Airlines which mentions use of Kubernetes. <https://www.linkedin.com/in/rupa-m-b90836309/>.

---

<sup>1</sup> Plaintiffs do not accuse the public clouds of Defendant, to the extent those services are provided by a cloud provider with a license to Plaintiffs’ patents that covers Defendant’s activities. As a specific example, Plaintiffs do not accuse Amazon managed services, *i.e.*, Amazon Elastic Kubernetes Service (Amazon EKS) and Amazon Elastic Container Service (Amazon ECS). Plaintiffs also does not accuse IBM managed services, *i.e.*, Red Hat Open Shift. Further, Plaintiffs do not accuse the public clouds of Defendants if those services are provided by a cloud provider with a license to Plaintiffs’ patents that covers Defendants’ activities. Plaintiffs will produce relevant license agreements in this litigation. Plaintiffs accuse Defendant private clouds that implement Kubernetes and non-licensed public clouds that Defendant uses to support Kubernetes for its systems and services. Plaintiffs will provide relevant license agreements for cloud providers in discovery, to the extent any such license agreements have not already been produced. To the extent any of these licenses are relevant to Defendant’s activities, Plaintiffs will meet and confer with Defendant about the impact of such license(s).

- Example of DevOps Engineer position at American Airlines which mentions use of Kubernetes. <https://www.linkedin.com/in/manidhar-a-555726169/>.
- Example of Java Developer position at American Airlines which mentions use of Kubernetes. <https://www.linkedin.com/in/suganya-koodalingam-8a0590102/>.
- Example of Senior Java Developer position at American Airlines which mentions use of Kubernetes. <https://www.linkedin.com/in/ganesh-kenda/>.

As another example, American has announced cloud migration of legacy technology and efforts to modernize its mainframes and servers. Source: <https://dxc.com/sg/en/insights/customer-stories/american-airlines-cloud-data-automation>. American continues to use private cloud for at least certain applications. Source: <https://www.techtarget.com/searchdatamanagement/feature/American-Airlines-lowers-data-management-costs-with-Intel> (“American Airlines’ initial target for cost optimization was Azure Data Lake, according to Vijay Premkumar, senior manager of public and *private cloud* at the airline.”) (emphasis added).

On information and belief, additional information confirms American uses Kubernetes technology.

**Q2: You've been leading digital transformation and modernization initiatives for a while, in very complex, demanding organizations. Looking back at your career, what are some of the challenges (organizational, business, and technical) that you encountered with modernization initiatives?**

**Jason:** In late 2018, we started our digital transformation by creating a tenancy in Oracle Cloud Infrastructure (OCI) and iterating through multiple proofs of technology, first proving out that our on-premises applications could not only run functionally, but also run workloads as fast or even faster in OCI. Shortly after that, a core group of infrastructure team members (led by our brilliant technical delivery manager, Vijay Krishnaswamy) began learning the fundamentals of Kubernetes and related open source software. Despite their team having success migrating their applications into OCI in 2020 and then redeploying their applications into Kubernetes a year later, we faced a larger challenge with the team supporting Ventana, the core application that supports the AAdvantage program. The stakes were huge: Ventana is one of American Airlines’ most important applications. Additionally, the team supporting it had not been directly involved in the journey forged by the other group.

When we merged both teams into one organization, encouraged team members to learn from and teach each other, and demonstrated sustained leadership support, we emerged with a crossfunctional team who achieved multiple certifications in both OCI and Kubernetes. They were then equipped with the skills, confidence, and support to move Ventana into not just OCI, but also Oracle Container Engine for Kubernetes (OKE) in one go.

Source: <https://blogs.oracle.com/cloud-infrastructure/post/five-questions-jason-maczura-american-airlines>.<sup>2</sup>

---

<sup>2</sup> Unless otherwise noted, all sources cited in this document were publicly accessible as of the date of the Amended Complaint.

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
[1.pre]. A computer system, comprising:	<p>To the extent this preamble is limiting, on information and belief, the American Count IX Systems and Services include a computer system.</p> <p>Kubernetes hosted on a server system, is a portable, extensible, open-source platform for managing containerized workloads and services.</p> <h2>Overview</h2> <p>Kubernetes is a portable, extensible, open source platform for managing containerized workloads and services, that facilitates both declarative configuration and automation. It has a large, rapidly growing ecosystem. Kubernetes services, support, and tools are widely available.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/overview/">https://kubernetes.io/docs/concepts/overview/</a>.</p> <p>Further, Kubernetes uses clusters, which are groups of nodes that host and run containerized applications as per their defined Deployments and Services.</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<h2>Cluster Architecture</h2> <p>The architectural concepts behind Kubernetes.</p> <p>A Kubernetes cluster consists of a control plane plus a set of worker machines, called nodes, that run containerized applications. Every cluster needs at least one worker node in order to run Pods.</p> <p>The worker node(s) host the Pods that are the components of the application workload. The control plane manages the worker nodes and the Pods in the cluster. In production environments, the control plane usually runs across multiple computers and a cluster usually runs multiple nodes, providing fault-tolerance and high availability.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/architecture/">https://kubernetes.io/docs/concepts/architecture/</a>.</p>

**U.S. Patent No. 8,352,584 (Claim 1)**

**Claim 1**

**Example American Count IX Systems and Services**

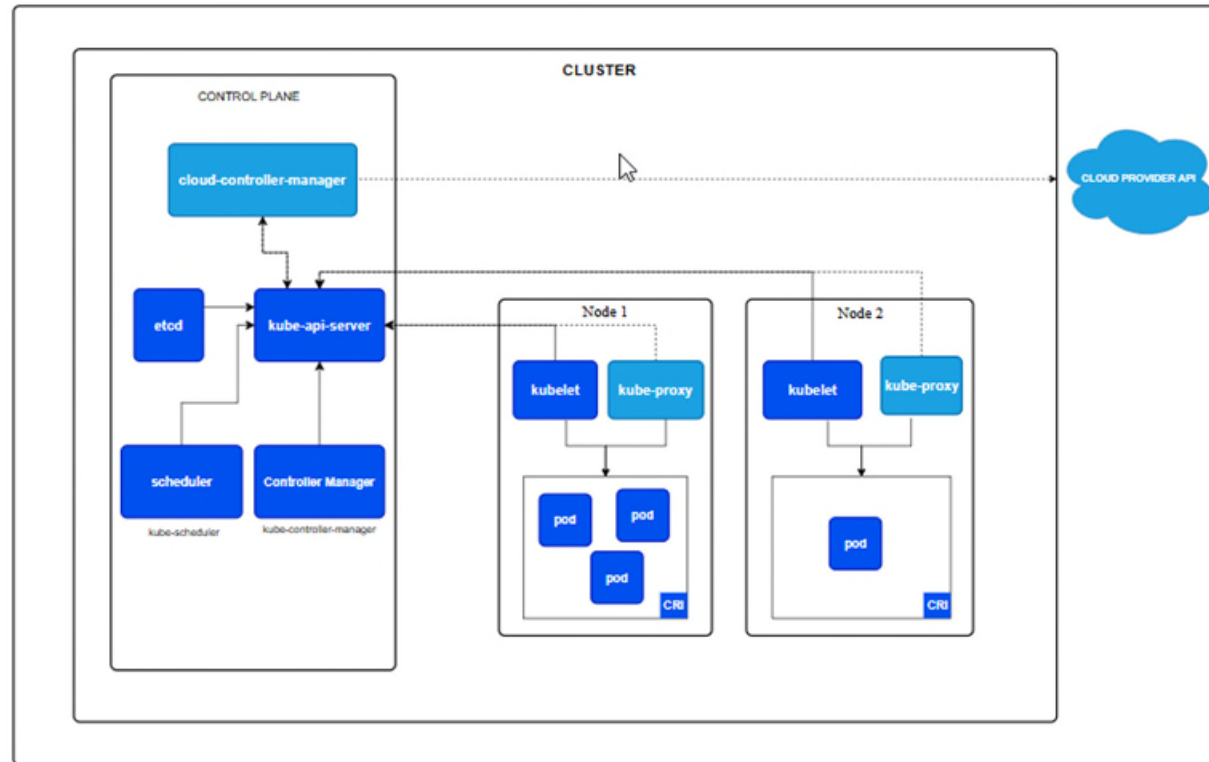


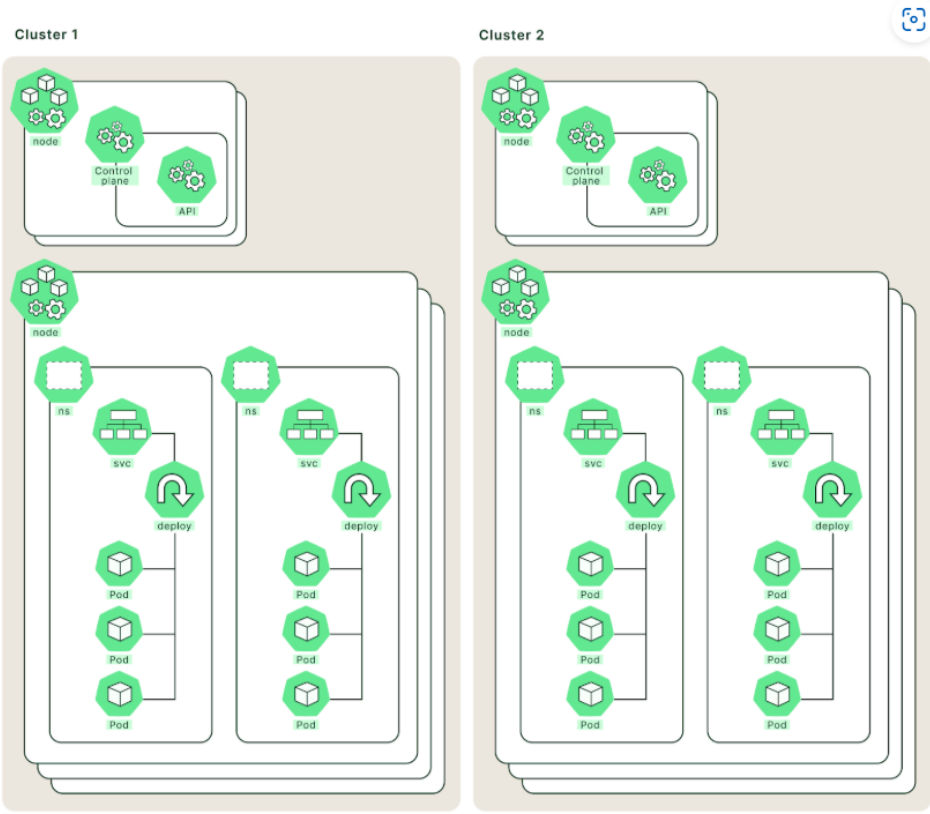
Figure 1. Kubernetes cluster components.

Source: <https://kubernetes.io/docs/concepts/architecture/>.

As shown below, Kubernetes supports multiple clusters, which can be hosted on a server system.

Kubernetes is designed to scale quickly, massively, and reliably. However, scaling in place — such as adding more nodes to a cluster — eventually reaches a point of diminishing returns. Kubernetes multi-cluster, which is multiple Kubernetes clusters operating as a single logical platform, helps solve this problem and enhances Kubernetes capabilities in many critical areas.

Source: <https://www.kubecost.com/kubernetes-multi-cloud/kubernetes-multi-cluster/>.


U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	 <p>Source: <a href="https://www.kubecost.com/kubernetes-multi-cloud/kubernetes-multi-cluster/">https://www.kubecost.com/kubernetes-multi-cloud/kubernetes-multi-cluster/</a>.</p>
[1.a] a private communications network linked to a public communications network;	<p>On information and belief, the American Count IX Systems and Services include a private communications network linked to a public communications network.</p> <p>For example, the Kubernetes system within the American accused network is configured to be linked through, for example, an Ingress, to an external network such as the Internet.</p>



U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p>For example, Kubernetes supports a Kubernetes cluster network to facilitate communication amongst nodes and/or Pods within the American accused network.</p> <h2>Terminology</h2> <p>For clarity, this guide defines the following terms:</p> <ul style="list-style-type: none"><li>• Node: A worker machine in Kubernetes, part of a cluster.</li><li>• Cluster: A set of Nodes that run containerized applications managed by Kubernetes. For this example, and in most common Kubernetes deployments, nodes in the cluster are not part of the public internet.</li><li>• Edge router: A router that enforces the firewall policy for your cluster. This could be a gateway managed by a cloud provider or a physical piece of hardware.</li><li>• Cluster network: A set of links, logical or physical, that facilitate communication within a cluster according to the Kubernetes <a href="#">networking model</a>.</li><li>• Service: A Kubernetes Service that identifies a set of Pods using <a href="#">label selectors</a>. Unless mentioned otherwise, Services are assumed to have virtual IPs only routable within the cluster network.</li></ul> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/ingress/">https://kubernetes.io/docs/concepts/services-networking/ingress/</a>.</p> <h2>The Kubernetes network model</h2>




U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<ul style="list-style-type: none"> <li>• The <i>pod network</i> (also called a cluster network) handles communication between pods. It ensures that (barring intentional network segmentation): <ul style="list-style-type: none"> <li>◦ All pods can communicate with all other pods, whether they are on the same <b>node</b> or on different nodes. Pods can communicate with each other directly, without the use of proxies or address translation (NAT).</li> </ul> <p>On Windows, this rule does not apply to host-network pods.</p> <li>◦ Agents on a node (such as system daemons, or kubelet) can communicate with all pods on that node.</li> </li></ul> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/">https://kubernetes.io/docs/concepts/services-networking/</a>.</p> <p>Further, in Kubernetes, Services provide a way for communication within the American accused network to an external network, such as the Internet. For example, the Ingress is configured to expose HTTP/HTTPS routes from outside the cluster.</p> <p><b>Ingress</b> exposes HTTP and HTTPS routes from outside the cluster to <b>services</b> within the cluster. Traffic routing is controlled by rules defined on the Ingress resource.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/ingress/">https://kubernetes.io/docs/concepts/services-networking/ingress/</a>.</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	 <pre> graph LR     client((client)) -.-&gt; Ingress-managed load balancer  Ingress[Ingress]     subgraph cluster         Ingress -- routing rule --&gt; Service[Service]         Service --&gt; Pod1[Pod]         Service --&gt; Pod2[Pod]     end </pre> <p>The diagram illustrates the Ingress architecture. A client (represented by a grey oval) connects to an Ingress controller (represented by a blue rectangle) within a cluster. The connection is managed by an "Ingress-managed load balancer" (indicated by a dashed arrow). The Ingress controller then routes traffic, using a "routing rule" (indicated by a solid arrow), to a Service (represented by a blue rectangle). The Service then distributes traffic to multiple Pods (represented by blue rectangles) within the cluster.</p> <p>Figure. Ingress</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/ingress/">https://kubernetes.io/docs/concepts/services-networking/ingress/</a>.</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p><b>Service</b></p> <p>Expose an application running in your cluster behind a single outward-facing endpoint, even when the workload is split across multiple backends.</p> <p>In Kubernetes, a Service is a method for exposing a network application that is running as one or more <u>Pods</u> in your cluster.</p> <p>A key aim of Services in Kubernetes is that you don't need to modify your existing application to use an unfamiliar service discovery mechanism. You can run code in Pods, whether this is a code designed for a cloud-native world, or an older app you've containerized. You use a Service to make that set of Pods available on the network so that clients can interact with it.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/service/">https://kubernetes.io/docs/concepts/services-networking/service/</a>.</p> <p>If your workload speaks HTTP, you might choose to use an <u>Ingress</u> to control how web traffic reaches that workload. Ingress is not a Service type, but it acts as the entry point for your cluster. An Ingress lets you consolidate your routing rules into a single resource, so that you can expose multiple components of your workload, running separately in your cluster, behind a single listener.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/service/">https://kubernetes.io/docs/concepts/services-networking/service/</a>.</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p>A Kubernetes system from within the American accused network includes the Gateway API, which provides Gateway resources. These resources can be defined and specified so that external traffic from the Internet can be routed to various Services.</p> <ul style="list-style-type: none"> <li>• The <a href="#">Gateway</a> API (or its predecessor, <a href="#">Ingress</a>) allows you to make Services accessible to clients that are outside the cluster.</li> </ul> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/">https://kubernetes.io/docs/concepts/services-networking/</a>.</p> <h2>Resource model</h2> <p>Gateway API has three stable API kinds:</p> <ul style="list-style-type: none"> <li>• <b>GatewayClass:</b> Defines a set of gateways with common configuration and managed by a controller that implements the class.</li> <li>• <b>Gateway:</b> Defines an instance of traffic handling infrastructure, such as cloud load balancer.</li> <li>• <b>HTTPRoute:</b> Defines HTTP-specific rules for mapping traffic from a Gateway listener to a representation of backend network endpoints. These endpoints are often represented as a <a href="#">Service</a>.</li> </ul> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/gateway/">https://kubernetes.io/docs/concepts/services-networking/gateway/</a>.</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	 <pre> graph LR     client([client]) -.-&gt; HTTP request  Gateway[Gateway]     Gateway --&gt; HTTPRoute[HTTPRoute]     HTTPRoute --&gt; Routing rule  Service[Service]     Service --&gt; Pod1[Pod]     Service --&gt; Pod2[Pod] </pre> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/gateway/">https://kubernetes.io/docs/concepts/services-networking/gateway/</a>.</p> <p>In this example, the request flow for a Gateway implemented as a reverse proxy is:</p> <ol style="list-style-type: none"> <li>1. The client starts to prepare an HTTP request for the URL  <code>http://www.example.com</code></li> </ol> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/gateway/">https://kubernetes.io/docs/concepts/services-networking/gateway/</a>.</p>
[1.b] a first cluster comprising a set of computing resources, including at least one hardware processor, in a first configuration, wherein the first cluster is communicatively linked to the private communications network;	<p>On information and belief, the American Count IX Systems and Services include a first cluster comprising a set of computing resources, including at least one hardware processor, in a first configuration, wherein the first cluster is communicatively linked to the private communications network.</p> <p>For example, Kubernetes clusters include at least one node, where a node is either a physical or virtual machine comprising a CPU or portion of CPU resources and memory, to run workloads. These clusters are connected to a private network, for example, the American accused network. On information and belief, the private network, for example, the American accused network(s), facilitates container-to-container, Pod-to-Pod, and/or Pod-to-Services communications across multiple clusters for sharing data and handling tasks.</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<h2>Cluster Architecture</h2> <p>The architectural concepts behind Kubernetes.</p> <p>A Kubernetes cluster consists of a control plane plus a set of worker machines, called nodes, that run containerized applications. Every cluster needs at least one worker node in order to run Pods.</p> <p>The worker node(s) host the Pods that are the components of the application workload. The control plane manages the worker nodes and the Pods in the cluster. In production environments, the control plane usually runs across multiple computers and a cluster usually runs multiple nodes, providing fault-tolerance and high availability.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/architecture/">https://kubernetes.io/docs/concepts/architecture/</a>.</p>



**U.S. Patent No. 8,352,584 (Claim 1)**

**Claim 1**

**Example American Count IX Systems and Services**

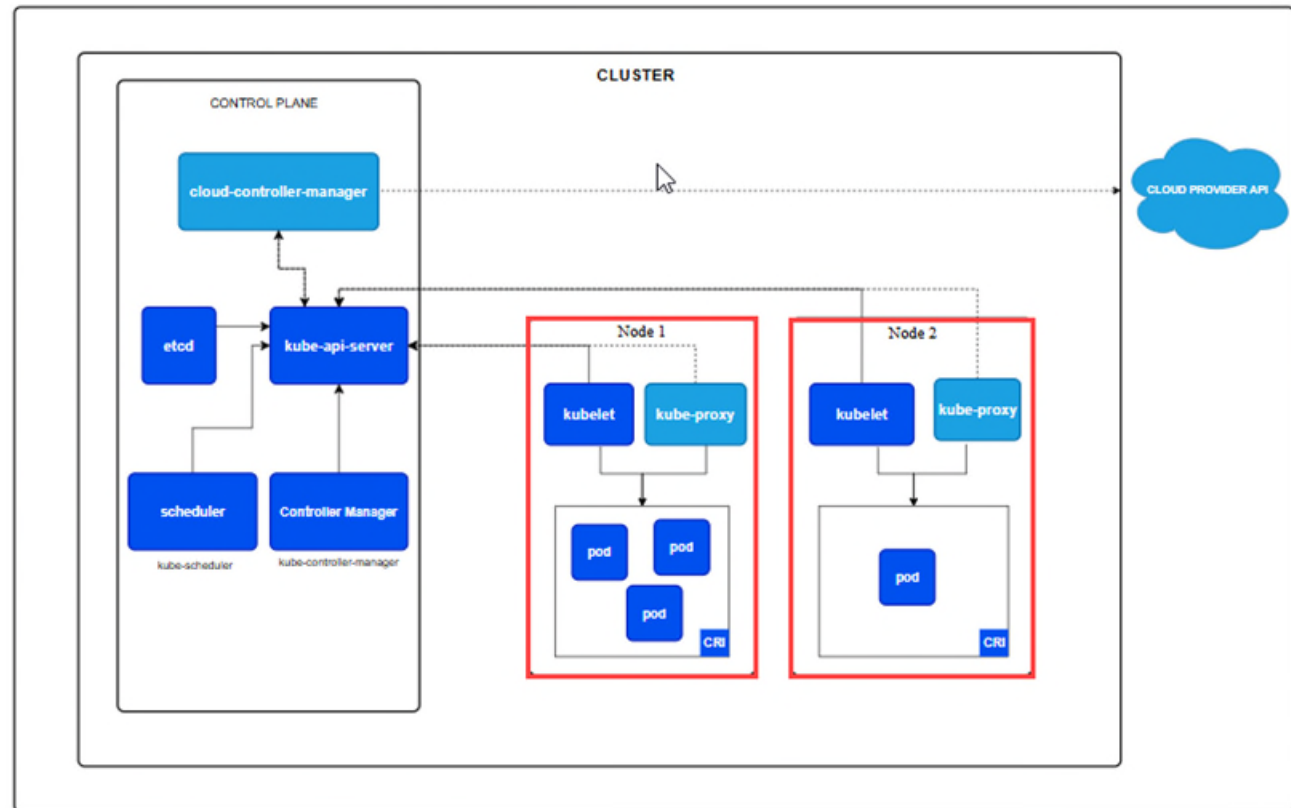


Figure 1. Kubernetes cluster components.

Source: <https://kubernetes.io/docs/concepts/architecture/>.


U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p>Networking is a central part of Kubernetes, but it can be challenging to understand exactly how it is expected to work. There are 4 distinct networking problems to address:</p> <ol style="list-style-type: none"> <li>1. Highly-coupled container-to-container communications: this is solved by <u>Pods</u> and <u>localhost</u> communications.</li> <li>2. Pod-to-Pod communications: this is the primary focus of this document.</li> <li>3. Pod-to-Service communications: this is covered by <u>Services</u>.</li> <li>4. External-to-Service communications: this is also covered by <u>Services</u>.</li> </ol> <p>Source: <a href="https://kubernetes.io/docs/concepts/cluster-administration/networking/">https://kubernetes.io/docs/concepts/cluster-administration/networking/</a>.</p> <h2>Nodes</h2> <p>Kubernetes runs your <u>workload</u> by placing containers into Pods to run on <i>Nodes</i>. A node may be a virtual or physical machine, depending on the cluster. Each node is managed by the control plane and contains the services necessary to run <u>Pods</u>.</p> <p>Typically you have several nodes in a cluster; in a learning or resource-limited environment, you might have only one node.</p> <p>The <u>components</u> on a node include the <u>kubelet</u>, a container runtime, and the <u>kube-proxy</u>.</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p>Node objects track information about the Node's resource capacity: for example, the amount of memory available and the number of CPUs. Nodes that self register report their capacity during registration. If you manually add a Node, then you need to set the node's capacity information when you add it.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/architecture/nodes/">https://kubernetes.io/docs/concepts/architecture/nodes/</a>.</p> <p>Kubernetes supports multi-cluster architectures, for example, such as cluster 1 and cluster 2 below, which can be hosted on a server system.</p> <p>Kubernetes is designed to scale quickly, massively, and reliably. However, scaling in place — such as adding more nodes to a cluster — eventually reaches a point of diminishing returns. Kubernetes multi-cluster, which is multiple Kubernetes clusters operating as a single logical platform, helps solve this problem and enhances Kubernetes capabilities in many critical areas.</p> <p>Source: <a href="https://www.kubecost.com/kubernetes-multi-cloud/kubernetes-multi-cluster/">https://www.kubecost.com/kubernetes-multi-cloud/kubernetes-multi-cluster/</a>.</p>

U.S. Patent No. 8,352,584 (Claim 1)


Claim 1

Example American Count IX Systems and Services



Source: <https://www.kubecost.com/kubernetes-multi-cloud/kubernetes-multi-cluster/>.


U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p><b>Kubernetes Multi-Cluster: Why and When To Use Them</b></p> <p>Application containerization has disrupted the way software applications have been built and deployed. Over the years, Kubernetes has stood out as one of the best platforms for container orchestration. It has helped many companies achieve scalability, resilience, portability, and better resource utilization in their products. However, managing Kubernetes is still complex. The first question which comes to mind is whether we should use a single cluster or a multi-cluster for Kubernetes. Although a single cluster is easy to set up and manage and provides the basic features of Kubernetes, it lacks the typical resilience and high availability Kubernetes is famous for. In many cases, a single cluster is not enough to manage the load efficiently across all components. As a result, we need more than one cluster for a better division of workload and resources, hence the need for a multi-cluster solution.</p> <p>Source: <a href="https://www.qovery.com/blog/kubernetes-multi-cluster-why-and-when-to-use-them/">https://www.qovery.com/blog/kubernetes-multi-cluster-why-and-when-to-use-them/</a>.</p> <p>On information and belief, in Kubernetes, Services provide a way for communication within the American accused network to an external network, such as the Internet. For example, the Ingress is configured to expose HTTP/HTTPS routes from outside the cluster.</p> <p>Ingress exposes HTTP and HTTPS routes from outside the cluster to services within the cluster. Traffic routing is controlled by rules defined on the Ingress resource.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/ingress/">https://kubernetes.io/docs/concepts/services-networking/ingress/</a>.</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	 <pre> graph LR     client([client]) -.-&gt; Ingress-managed load balancer  Ingress[Ingress]     subgraph cluster         Ingress -- routing rule --&gt; Service[Service]         Service --&gt; Pod1[Pod]         Service --&gt; Pod2[Pod]     end             </pre> <p>The diagram illustrates the Ingress architecture. A client (represented by a grey oval) connects to an Ingress controller (represented by a blue rectangle) within a cluster. The connection is managed by an "Ingress-managed load balancer" (indicated by a dashed arrow). The Ingress controller then routes traffic, using a "routing rule" (indicated by a solid arrow), to a Service (represented by a blue rectangle). The Service then distributes traffic to multiple Pods (represented by blue rectangles) within the cluster.</p> <p>Figure. Ingress</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/ingress/">https://kubernetes.io/docs/concepts/services-networking/ingress/</a>.</p>



U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p><b>Service</b></p> <p>Expose an application running in your cluster behind a single outward-facing endpoint, even when the workload is split across multiple backends.</p> <p>In Kubernetes, a Service is a method for exposing a network application that is running as one or more <u>Pods</u> in your cluster.</p> <p>A key aim of Services in Kubernetes is that you don't need to modify your existing application to use an unfamiliar service discovery mechanism. You can run code in Pods, whether this is a code designed for a cloud-native world, or an older app you've containerized. You use a Service to make that set of Pods available on the network so that clients can interact with it.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/service/">https://kubernetes.io/docs/concepts/services-networking/service/</a>.</p> <p>If your workload speaks HTTP, you might choose to use an <u>Ingress</u> to control how web traffic reaches that workload. Ingress is not a Service type, but it acts as the entry point for your cluster. An Ingress lets you consolidate your routing rules into a single resource, so that you can expose multiple components of your workload, running separately in your cluster, behind a single listener.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/service/">https://kubernetes.io/docs/concepts/services-networking/service/</a>.</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p>On information and belief, a Kubernetes system from within the American accused network includes the Gateway API, which provides Gateway resources. These resources can be defined and specified so that external traffic from the Internet can be routed to various Services.</p> <ul style="list-style-type: none"> <li>• The <a href="#">Gateway</a> API (or its predecessor, <a href="#">Ingress</a>) allows you to make Services accessible to clients that are outside the cluster.</li> </ul> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/">https://kubernetes.io/docs/concepts/services-networking/</a>.</p> <h2>Resource model</h2> <p>Gateway API has three stable API kinds:</p> <ul style="list-style-type: none"> <li>• <b>GatewayClass:</b> Defines a set of gateways with common configuration and managed by a controller that implements the class.</li> <li>• <b>Gateway:</b> Defines an instance of traffic handling infrastructure, such as cloud load balancer.</li> <li>• <b>HTTPRoute:</b> Defines HTTP-specific rules for mapping traffic from a Gateway listener to a representation of backend network endpoints. These endpoints are often represented as a <a href="#">Service</a>.</li> </ul> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/gateway/">https://kubernetes.io/docs/concepts/services-networking/gateway/</a>.</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	 <pre> graph LR     client([client]) -.-&gt; HTTP request  Gateway[Gateway]     Gateway --&gt; HTTPRoute[HTTPRoute]     HTTPRoute --&gt; Routing rule  Service[Service]     Service --&gt; Pod1[Pod]     Service --&gt; Pod2[Pod] </pre> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/gateway/">https://kubernetes.io/docs/concepts/services-networking/gateway/</a>.</p> <p>In this example, the request flow for a Gateway implemented as a reverse proxy is:</p> <ol style="list-style-type: none"> <li>1. The client starts to prepare an HTTP request for the URL  <code>http://www.example.com</code></li> </ol> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/gateway/">https://kubernetes.io/docs/concepts/services-networking/gateway/</a>.</p>
[1.c] a second cluster comprising a set of computing resources, including at least one hardware processor, in a second configuration, wherein the second cluster is communicatively linked to the private communications network; and	<p>On information and belief, the American Count IX Systems and Services include a second cluster comprising a set of computing resources, including at least one hardware processor, in a second configuration, wherein the second cluster is communicatively linked to the private communications network.</p> <p>For example, Kubernetes clusters include at least one node, where a node is either a physical or virtual machine comprising a CPU or portion of CPU resources and memory, to run workloads. These clusters are connected to a private network, for example, the American accused network. On information and belief, the private network, for example, the American accused network, facilitates container-to-container, Pod-to-Pod, and/or Pod-to-Services communications across multiple clusters for sharing data and handling tasks.</p>


U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<h2>Cluster Architecture</h2> <p>The architectural concepts behind Kubernetes.</p> <p>A Kubernetes cluster consists of a control plane plus a set of worker machines, called nodes, that run containerized applications. Every cluster needs at least one worker node in order to run Pods.</p> <p>The worker node(s) host the Pods that are the components of the application workload. The control plane manages the worker nodes and the Pods in the cluster. In production environments, the control plane usually runs across multiple computers and a cluster usually runs multiple nodes, providing fault-tolerance and high availability.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/architecture/">https://kubernetes.io/docs/concepts/architecture/</a>.</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p><b>Nodes</b></p> <p>Kubernetes runs your workload by placing containers into Pods to run on <i>Nodes</i>. A node may be a virtual or physical machine, depending on the cluster. Each node is managed by the control plane and contains the services necessary to run Pods.</p> <p>Typically you have several nodes in a cluster; in a learning or resource-limited environment, you might have only one node.</p> <p>The components on a node include the kubelet, a container runtime, and the kube-proxy.</p> <p>Node objects track information about the Node's resource capacity: for example, the amount of memory available and the number of CPUs. Nodes that self register report their capacity during registration. If you manually add a Node, then you need to set the node's capacity information when you add it.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/architecture/nodes/">https://kubernetes.io/docs/concepts/architecture/nodes/</a>.</p> <p>Kubernetes supports multi-cluster architectures, for example, cluster 1 and cluster 2, which can be hosted on a server system.</p> <p>Kubernetes is designed to scale quickly, massively, and reliably. However, scaling in place — such as adding more nodes to a cluster — eventually reaches a point of diminishing returns. Kubernetes multi-cluster, which is multiple Kubernetes clusters operating as a single logical platform, helps solve this problem and enhances Kubernetes capabilities in many critical areas.</p> <p>Source: <a href="https://www.kubecost.com/kubernetes-multi-cloud/kubernetes-multi-cluster/">https://www.kubecost.com/kubernetes-multi-cloud/kubernetes-multi-cluster/</a>.</p>

U.S. Patent No. 8,352,584 (Claim 1)

Claim 1

Example American Count IX Systems and Services



Source: <https://www.kubecost.com/kubernetes-multi-cloud/kubernetes-multi-cluster/>.




U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p><b>Kubernetes Multi-Cluster: Why and When To Use Them</b></p> <p>Application containerization has disrupted the way software applications have been built and deployed. Over the years, Kubernetes has stood out as one of the best platforms for container orchestration. It has helped many companies achieve scalability, resilience, portability, and better resource utilization in their products. However, managing Kubernetes is still complex. The first question which comes to mind is whether we should use a single cluster or a multi-cluster for Kubernetes. Although a single cluster is easy to set up and manage and provides the basic features of Kubernetes, it lacks the typical resilience and high availability Kubernetes is famous for. In many cases, a single cluster is not enough to manage the load efficiently across all components. As a result, we need more than one cluster for a better division of workload and resources, hence the need for a multi-cluster solution.</p> <p>Source: <a href="https://www.qovery.com/blog/kubernetes-multi-cluster-why-and-when-to-use-them/">https://www.qovery.com/blog/kubernetes-multi-cluster-why-and-when-to-use-them/</a>.</p> <p>On information and belief, in Kubernetes, Services provide a way for communication within the American accused network to an external network, such as the Internet. For example, the Ingress is configured to expose HTTP/HTTPS routes from outside the cluster.</p> <p>Ingress exposes HTTP and HTTPS routes from outside the cluster to services within the cluster. Traffic routing is controlled by rules defined on the Ingress resource.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/ingress/">https://kubernetes.io/docs/concepts/services-networking/ingress/</a>.</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<div><pre>graph LR; client((client)) -.-&gt; Ingress-managed load balancer  Ingress[Ingress]; subgraph cluster; Ingress -- routing rule --&gt; Service[Service]; Service --&gt; Pod1[Pod]; Service --&gt; Pod2[Pod]; end</pre><p>The diagram illustrates the Ingress architecture. A client (represented by a grey oval) connects to an Ingress controller (represented by a blue rectangle) within a cluster. The connection is managed by an 'Ingress-managed load balancer' (indicated by a dashed arrow). The Ingress controller then routes traffic to a Service (represented by a blue rectangle) using a 'routing rule' (indicated by a solid arrow). The Service then routes traffic to multiple Pods (represented by blue rectangles) within the cluster (indicated by curved arrows).</p></div> <p>Figure. Ingress</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/ingress/">https://kubernetes.io/docs/concepts/services-networking/ingress/</a>.</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p><b>Service</b></p> <p>Expose an application running in your cluster behind a single outward-facing endpoint, even when the workload is split across multiple backends.</p> <p>In Kubernetes, a Service is a method for exposing a network application that is running as one or more <u>Pods</u> in your cluster.</p> <p>A key aim of Services in Kubernetes is that you don't need to modify your existing application to use an unfamiliar service discovery mechanism. You can run code in Pods, whether this is a code designed for a cloud-native world, or an older app you've containerized. You use a Service to make that set of Pods available on the network so that clients can interact with it.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/service/">https://kubernetes.io/docs/concepts/services-networking/service/</a>.</p> <p>If your workload speaks HTTP, you might choose to use an <u>Ingress</u> to control how web traffic reaches that workload. Ingress is not a Service type, but it acts as the entry point for your cluster. An Ingress lets you consolidate your routing rules into a single resource, so that you can expose multiple components of your workload, running separately in your cluster, behind a single listener.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/service/">https://kubernetes.io/docs/concepts/services-networking/service/</a>.</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p>On information and belief, a Kubernetes system from within the American accused network includes the Gateway API, which provides Gateway resources. These resources can be defined and specified so that external traffic from the Internet.</p> <ul style="list-style-type: none"> <li>• The <a href="#">Gateway</a> API (or its predecessor, <a href="#">Ingress</a>) allows you to make Services accessible to clients that are outside the cluster.</li> </ul> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/">https://kubernetes.io/docs/concepts/services-networking/</a>.</p> <h2>Resource model</h2> <p>Gateway API has three stable API kinds:</p> <ul style="list-style-type: none"> <li>• <b>GatewayClass:</b> Defines a set of gateways with common configuration and managed by a controller that implements the class.</li> <li>• <b>Gateway:</b> Defines an instance of traffic handling infrastructure, such as cloud load balancer.</li> <li>• <b>HTTPRoute:</b> Defines HTTP-specific rules for mapping traffic from a Gateway listener to a representation of backend network endpoints. These endpoints are often represented as a <a href="#">Service</a>.</li> </ul> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/gateway/">https://kubernetes.io/docs/concepts/services-networking/gateway/</a>.</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	 <pre> graph LR     client([client]) -.-&gt; HTTP request  Gateway[Gateway]     Gateway --&gt; HTTPRoute[HTTPRoute]     HTTPRoute -- "Routing rule" --&gt; Service[Service]     Service --&gt; Pod1[Pod]     Service --&gt; Pod2[Pod] </pre> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/gateway/">https://kubernetes.io/docs/concepts/services-networking/gateway/</a>.</p> <p>In this example, the request flow for a Gateway implemented as a reverse proxy is:</p> <ol style="list-style-type: none"> <li>1. The client starts to prepare an HTTP request for the URL  <code>http://www.example.com</code></li> </ol> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/gateway/">https://kubernetes.io/docs/concepts/services-networking/gateway/</a>.</p>
[1.d] a monitoring system to monitor operations of the first cluster and the second cluster for communications problems;	<p>On information and belief, the American Count IX Systems and Services include a monitoring system to monitor operations of the first cluster and the second cluster for communications problems.</p> <p>For example, based on information and belief, Kubernetes supports the Kubernetes API to manage and monitor Kubernetes clusters, for metrics such as cluster health, performance, and/or communication issues.</p> <p>Furthermore, Kubernetes works with full metrics pipeline and monitoring solutions that provide monitoring of operations within Kubernetes clusters, including connectivity problems.</p>

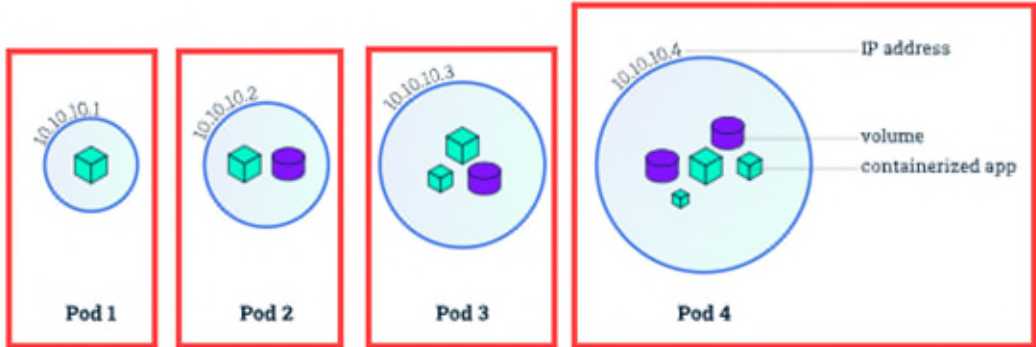
U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p>A full metrics pipeline gives you access to richer metrics. Kubernetes can respond to these metrics by automatically scaling or adapting the cluster based on its current state, using mechanisms such as the Horizontal Pod Autoscaler. The monitoring pipeline fetches metrics from the kubelet and then exposes them to Kubernetes via an adapter by implementing either the <code>custom.metrics.k8s.io</code> or <code>external.metrics.k8s.io</code> API.</p> <p>Kubernetes is designed to work with <a href="#">OpenMetrics</a>, which is one of the <a href="#">CNCF Observability and Analysis - Monitoring Projects</a>, built upon and carefully extending <a href="#">Prometheus exposition format</a> in almost 100% backwards-compatible ways.</p> <p>Source: <a href="https://kubernetes.io/docs/tasks/debug/debug-cluster/resource-usage-monitoring/">https://kubernetes.io/docs/tasks/debug/debug-cluster/resource-usage-monitoring/</a>.</p> <h2>What is Kubernetes monitoring?</h2> <p>Simply put, Kubernetes monitoring is the practice of tracking the status of all components of a Kubernetes environment. Because there are <a href="#">many pieces inside Kubernetes</a>, Kubernetes monitoring actually entails monitoring many distinct things, such as:</p> <ul style="list-style-type: none"> <li>• The kube-system workloads.</li> <li>• Cluster information using the Kubernetes API.</li> <li>• Applications interactions with Kubernetes by monitoring apps bottom-up.</li> </ul> <p>Source: <a href="https://www.groundcover.com/kubernetes-monitoring">https://www.groundcover.com/kubernetes-monitoring</a>.</p>



U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p>Kubernetes monitoring includes providing alerts to problems occurring within a cluster, which enables troubleshooting and addressing other potential issues.</p> <p>By collecting Kubernetes data, you'll get viable information regarding your Kubernetes cluster health, that can help you perform <b>Kubernetes troubleshooting</b> and manage issues like <b>unexpected container termination</b>. You can also leverage the data for proactive decisions such as adjusting <b>rate limits</b>.</p> <p>Source: <a href="https://www.groundcover.com/kubernetes-monitoring">https://www.groundcover.com/kubernetes-monitoring</a>.</p> <p>Furthermore, Kubernetes includes various metrics tracking tools, such as tools for monitoring node health across multiple clusters.</p> <h2>Monitor Node Health</h2> <p><i>Node Problem Detector</i> is a daemon for monitoring and reporting about a node's health. You can run Node Problem Detector as a <b>DaemonSet</b> or as a standalone daemon. Node Problem Detector collects information about node problems from various daemons and reports these conditions to the API server as Node <b>Conditions</b> or as <b>Events</b>.</p> <p>Source: <a href="https://kubernetes.io/docs/tasks/debug/debug-cluster/monitor-node-health/">https://kubernetes.io/docs/tasks/debug/debug-cluster/monitor-node-health/</a>.</p> <p>Kubernetes alerting provides for the identification of operational and connectivity problems, including the identification of clusters associated with the problem.</p>

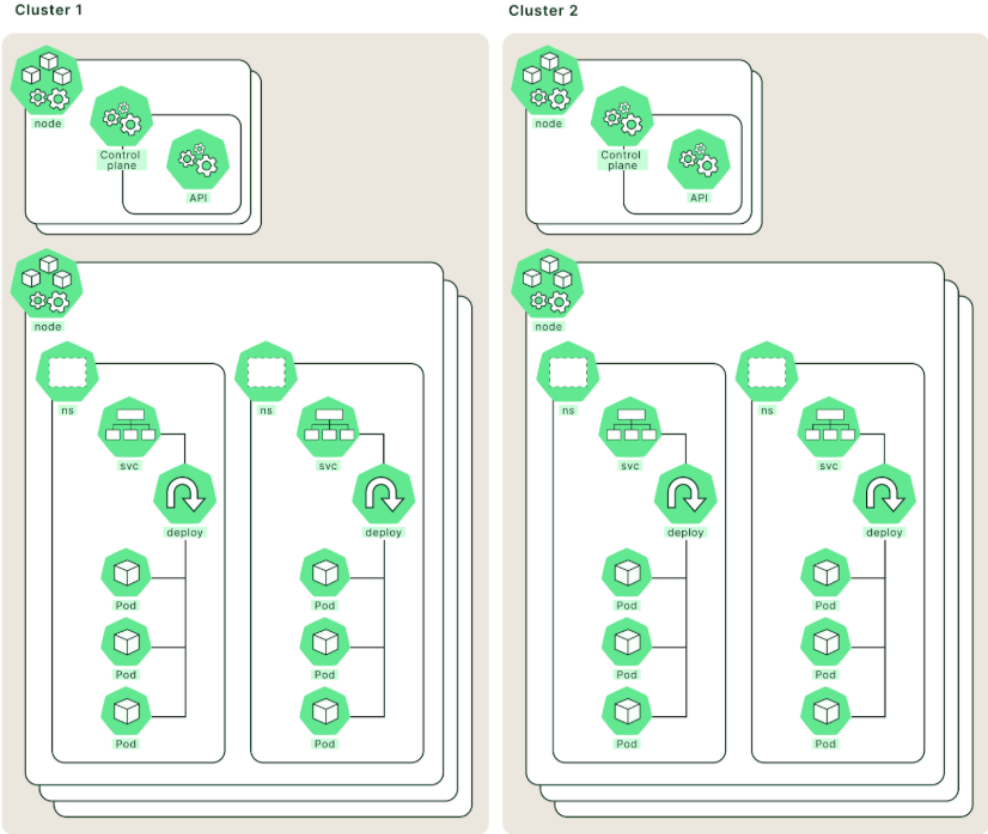
U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p>Kubernetes alerting is the practice of generating notifications for events or trends in Kubernetes that require admins' attention. Examples of such events and trends include:</p> <ul style="list-style-type: none"> <li>• A node that has failed.</li> <li>• A Pod that is stuck in the pending state.</li> <li>• A container or Pod that is consuming a high level of resources relative to normal consumption trends.</li> <li>• High latency rates for communication between Kubernetes cluster components (such as between kubelet and control plane nodes).</li> </ul> <p>Source: <a href="https://www.groundcover.com/kubernetes-monitoring/kubernetes-alerting">https://www.groundcover.com/kubernetes-monitoring/kubernetes-alerting</a>.</p>
[1.e] wherein the first configuration differs from the second configuration;	<p>On information and belief, the American Count IX Systems and Services include a system where the first configuration differs from the second configuration.</p> <p>For example, Kubernetes pods distributed across different nodes in different clusters can be configured to perform different tasks. Each node contains services necessary to run a pod, and each pod runs its own instance of a given application container or set of application containers, resulting in different configurations occurring at each node and cluster.</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<h2>Pods</h2> <p><i>Pods</i> are the smallest deployable units of computing that you can create and manage in Kubernetes.</p> <p>A <i>Pod</i> (as in a pod of whales or pea pod) is a group of one or more <u>containers</u>, with shared storage and network resources, and a specification for how to run the containers. A Pod's contents are always co-located and co-scheduled, and run in a shared context. A Pod models an application-specific "logical host": it contains one or more application containers which are relatively tightly coupled. In non-cloud contexts, applications executed on the same physical or virtual machine are analogous to cloud applications executed on the same logical host.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/workloads/pods/">https://kubernetes.io/docs/concepts/workloads/pods/</a>.</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p><b>Pods overview</b></p>  <p>Source: <a href="https://kubernetes.io/docs/tutorials/kubernetes-basics/explore/explore-intro/">https://kubernetes.io/docs/tutorials/kubernetes-basics/explore/explore-intro/</a>.</p> <p><b>Workload resources for managing pods</b> <a href="#">↗</a></p> <p>Usually you don't need to create Pods directly, even singleton Pods. Instead, create them using workload resources such as <u>Deployment</u> or <u>Job</u>. If your Pods need to track state, consider the <u>StatefulSet</u> resource.</p> <p>Each Pod is meant to run a single instance of a given application. If you want to scale your application horizontally (to provide more overall resources by running more instances), you should use multiple Pods, one for each instance. In Kubernetes, this is typically referred to as <i>replication</i>. Replicated Pods are usually created and managed as a group by a workload resource and its controller.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/workloads/pods/">https://kubernetes.io/docs/concepts/workloads/pods/</a>.</p>

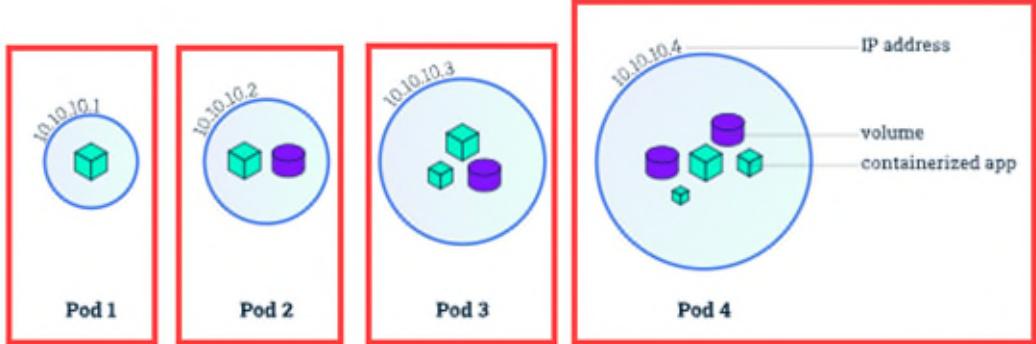
U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p>For example, components of the first cluster architecture differ from the second cluster architecture based on worker node configurations, the running containerized applications within Pods, networking, and other features part of Kubernetes.</p> <h2>Cluster Architecture</h2> <p>The architectural concepts behind Kubernetes.</p> <p>A Kubernetes cluster consists of a control plane plus a set of worker machines, called nodes, that run containerized applications. Every cluster needs at least one worker node in order to run Pods.</p> <p>The worker node(s) host the Pods that are the components of the application workload. The control plane manages the worker nodes and the Pods in the cluster. In production environments, the control plane usually runs across multiple computers and a cluster usually runs multiple nodes, providing fault-tolerance and high availability.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/architecture/">https://kubernetes.io/docs/concepts/architecture/</a>.</p>



U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	 <p>Source: <a href="https://www.kubecost.com/kubernetes-multi-cloud/kubernetes-multi-cluster/">https://www.kubecost.com/kubernetes-multi-cloud/kubernetes-multi-cluster/</a>.</p>
[1.f] wherein the first configuration provides a first computing environment to perform a first client task and the second configuration	<p>On information and belief, the American Count IX Systems and Services include a system where the first configuration provides a first computing environment to perform a first client task and the second configuration provides a second computing environment to perform a second client task.</p> <p>For example, Kubernetes pods distributed across different nodes in different clusters can be configured to perform different tasks. Each node contains services necessary to run a pod, and each pod runs its</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
provides a second computing environment to perform a second client task; and	<p>own instance of a given application container or set of application containers, resulting in different configurations occurring at each node and cluster.</p> <h2>Pods</h2> <p><i>Pods</i> are the smallest deployable units of computing that you can create and manage in Kubernetes.</p> <p>A <i>Pod</i> (as in a pod of whales or pea pod) is a group of one or more <u>containers</u>, with shared storage and network resources, and a specification for how to run the containers. A Pod's contents are always co-located and co-scheduled, and run in a shared context. A Pod models an application-specific "logical host": it contains one or more application containers which are relatively tightly coupled. In non-cloud contexts, applications executed on the same physical or virtual machine are analogous to cloud applications executed on the same logical host.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/workloads/pods/">https://kubernetes.io/docs/concepts/workloads/pods/</a>.</p>




U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p><b>Pods overview</b></p>  <p>Source: <a href="https://kubernetes.io/docs/tutorials/kubernetes-basics/explore/explore-intro/">https://kubernetes.io/docs/tutorials/kubernetes-basics/explore/explore-intro/</a>.</p> <p><b>Workload resources for managing pods</b> <a href="#">↗</a></p> <p>Usually you don't need to create Pods directly, even singleton Pods. Instead, create them using workload resources such as <u>Deployment</u> or <u>Job</u>. If your Pods need to track state, consider the <u>StatefulSet</u> resource.</p> <p>Each Pod is meant to run a single instance of a given application. If you want to scale your application horizontally (to provide more overall resources by running more instances), you should use multiple Pods, one for each instance. In Kubernetes, this is typically referred to as <i>replication</i>. Replicated Pods are usually created and managed as a group by a workload resource and its <u>controller</u>.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/workloads/pods/">https://kubernetes.io/docs/concepts/workloads/pods/</a>.</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p>For example, components of the first cluster architecture differ from the second cluster architecture based on worker node configurations, the running containerized applications within Pods, networking, and other features part of Kubernetes.</p> <h2>Cluster Architecture</h2> <p>The architectural concepts behind Kubernetes.</p> <p>A Kubernetes cluster consists of a control plane plus a set of worker machines, called nodes, that run containerized applications. Every cluster needs at least one worker node in order to run Pods.</p> <p>The worker node(s) host the Pods that are the components of the application workload. The control plane manages the worker nodes and the Pods in the cluster. In production environments, the control plane usually runs across multiple computers and a cluster usually runs multiple nodes, providing fault-tolerance and high availability.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/architecture/">https://kubernetes.io/docs/concepts/architecture/</a>.</p>

**U.S. Patent No. 8,352,584 (Claim 1)**

**Claim 1**

**Example American Count IX Systems and Services**



Source: <https://www.kubecost.com/kubernetes-multi-cloud/kubernetes-multi-cluster/>.

As a further example, multiple cluster configurations can differ based on different environmental access, location of the cluster, performance, and other considerations as mentioned below.

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p>There are many reasons to want to run multiple clusters, including but not limited to:</p> <ul style="list-style-type: none"> <li>• Location <ul style="list-style-type: none"> <li>• Latency: it can be necessary to deploy the application as close to the customers as possible.</li> <li>• Jurisdiction: it can be mandated to keep user data in-country.</li> <li>• Data gravity: data already exists in one provider, but it can be decided to run the application in another environment.</li> </ul> </li> <li>• Isolation <ul style="list-style-type: none"> <li>• Environment (e.g. dev, test, prod)</li> <li>• Performance isolation: a workload may consume too many resources, at the expense of other workloads.</li> <li>• Security isolation: sensitive data or untrusted code must be isolated in their own environments.</li> <li>• Organizational isolation: teams may have different management domains.</li> <li>• Cost isolation: multitenancy can greatly complexify billing management for different teams.</li> </ul> </li> <li>• Reliability <ul style="list-style-type: none"> <li>• Blast radius: an infrastructure or application incident in one cluster must not impact the whole system.</li> <li>• Infrastructure diversity: an underlying zone, region, or provider outage does not bring down the whole system.</li> <li>• Scale: the application is too big to fit in a single cluster.</li> <li>• Upgrade scope: some parts of the application may require an infrastructure upgrade, that may impact other parts of the application. Having multiple clusters can also avoid the need for in-place cluster upgrades.</li> </ul> </li> </ul> <p>Source: <a href="https://multicluster.sigs.k8s.io/">https://multicluster.sigs.k8s.io/</a>.</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p>For example, Kubernetes implements network policies for controlling the traffic at the IP address or port level within the cluster, and between Pods and the outside world. Network policies limit access by only allowing specific IP ranges to request.</p> <h2>Network Policies</h2> <p>If you want to control traffic flow at the IP address or port level (OSI layer 3 or 4), NetworkPolicies allow you to specify rules for traffic flow within your cluster, and also between Pods and the outside world. Your cluster must use a network plugin that supports NetworkPolicy enforcement.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/network-policies/">https://kubernetes.io/docs/concepts/services-networking/network-policies/</a>.</p> <p>The entities that a Pod can communicate with are identified through a combination of the following three identifiers:</p> <ol style="list-style-type: none"><li>1. Other pods that are allowed (exception: a pod cannot block access to itself)</li><li>2. Namespaces that are allowed</li><li>3. IP blocks (exception: traffic to and from the node where a Pod is running is always allowed, regardless of the IP address of the Pod or the node)</li></ol> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/network-policies/">https://kubernetes.io/docs/concepts/services-networking/network-policies/</a>.</p>

```
service/networking/networkpolicy.yaml

apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
  name: test-network-policy
  namespace: default
spec:
  podSelector:
    matchLabels:
      role: db
  policyTypes:
  - Ingress
  - Egress
  ingress:
  - from:
    - ipBlock:
        cidr: 172.17.0.0/16
        except:
        - 172.17.1.0/24
    - namespaceSelector:
        matchLabels:
          project: myproject
    - podSelector:
        matchLabels:
          role: frontend
  ports:
  - protocol: TCP
    port: 6379
  egress:
  - to:
    - ipBlock:
        cidr: 10.0.0.0/24
  ports:
  - protocol: TCP
    port: 5978
```

Source: <https://kubernetes.io/docs/concepts/services-networking/network-policies/>.

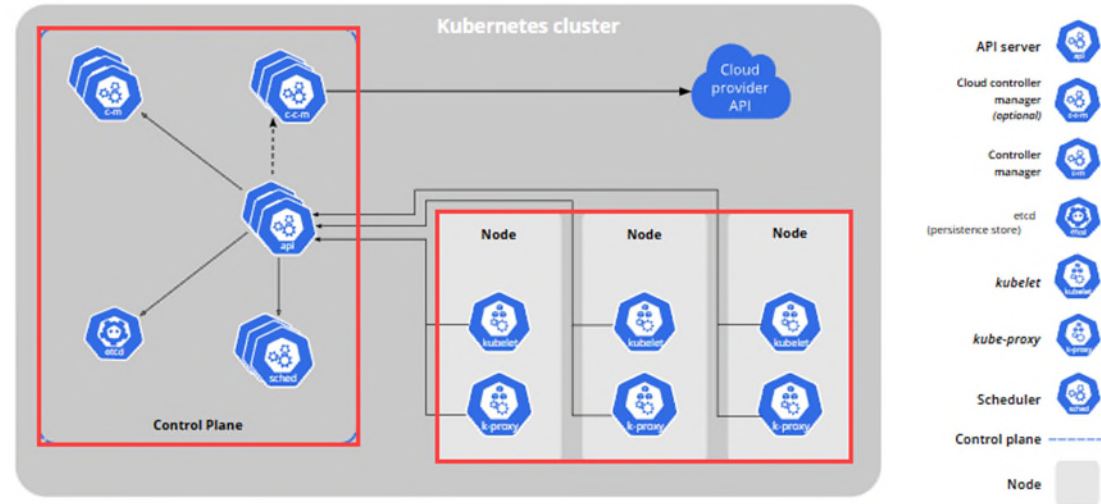
U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p>So, the example NetworkPolicy:</p> <ol style="list-style-type: none"> <li>1. isolates <code>role=db</code> pods in the <code>default</code> namespace for both ingress and egress traffic (if they weren't already isolated)</li> <li>2. (Ingress rules) allows connections to all pods in the <code>default</code> namespace with the label <code>role=db</code> on TCP port 6379 from: <ul style="list-style-type: none"> <li>◦ any pod in the <code>default</code> namespace with the label <code>role=frontend</code></li> <li>◦ any pod in a namespace with the label <code>project=myproject</code></li> <li>◦ IP addresses in the ranges <code>172.17.0.0 – 172.17.0.255</code> and <code>172.17.2.0 – 172.17.255.255</code> (ie, all of <code>172.17.0.0/16</code> except <code>172.17.1.0/24</code> )</li> </ul> </li> <li>3. (Egress rules) allows connections from any pod in the <code>default</code> namespace with the label <code>role=db</code> to CIDR <code>10.0.0.0/24</code> on TCP port 5978</li> </ol> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/network-policies/">https://kubernetes.io/docs/concepts/services-networking/network-policies/</a>.</p>
[1.g] wherein the computing resources comprise processing nodes, data storage shared by the processing nodes, and at least one communications network to link the processing nodes to each other and to the data storage;	<p>On information and belief, the American Count IX Systems and Services include a system where the computing resources comprise processing nodes, data storage shared by the processing nodes, and at least one communications network to link the processing nodes to each other and to the data storage.</p> <p>For example, organized groups of resources such as CPU and memory are accessed by nodes of Kubernetes worker machines. Nodes are managed by the Kubernetes control plane, where a pod and its containers run on a node.</p>



**U.S. Patent No. 8,352,584 (Claim 1)**

**Claim 1**

**Example American Count IX Systems and Services**



The components of a Kubernetes cluster

Source: <https://kubernetes.io/docs/concepts/overview/components/>.

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<h2>Resource Management for Pods and Containers</h2> <p>When you specify a <u>Pod</u>, you can optionally specify how much of each resource a <u>container</u> needs. The most common resources to specify are CPU and memory (RAM); there are others.</p> <p>When you specify the resource <i>request</i> for containers in a Pod, the <u>kube-scheduler</u> uses this information to decide which node to place the Pod on. When you specify a resource <i>limit</i> for a container, the <u>kubelet</u> enforces those limits so that the running container is not allowed to use more of that resource than the limit you set. The kubelet also reserves at least the <i>request</i> amount of that system resource specifically for that container to use.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/">https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/</a>.</p>


U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p><b>Core Components</b></p> <p>A Kubernetes cluster consists of a control plane and one or more worker nodes. Here's a brief overview of the main components:</p> <p><b>Control Plane Components</b></p> <p>Manage the overall state of the cluster:</p> <p><b>kube-apiserver</b></p> <p>The core component server that exposes the Kubernetes HTTP API</p> <p><b>etcd</b></p> <p>Consistent and highly-available key value store for all API server data</p> <p><b>kube-scheduler</b></p> <p>Looks for Pods not yet bound to a node, and assigns each Pod to a suitable node.</p> <p><b>kube-controller-manager</b></p> <p>Runs <u>controllers</u> to implement Kubernetes API behavior.</p> <p><b>cloud-controller-manager (optional)</b></p> <p>Integrates with underlying cloud provider(s).</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/overview/components/">https://kubernetes.io/docs/concepts/overview/components/</a>.</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p><b>Node Components</b></p> <p>Run on every node, maintaining running pods and providing the Kubernetes runtime environment:</p> <p><b>kubelet</b></p> <p>Ensures that Pods are running, including their containers.</p> <p><b>kube-proxy (optional)</b></p> <p>Maintains network rules on nodes to implement <u>Services</u>.</p> <p><b>Container runtime</b></p> <p>Software responsible for running containers. Read <a href="#">Container Runtimes</a> to learn more.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/overview/components/">https://kubernetes.io/docs/concepts/overview/components/</a>.</p> <p>Furthermore, Kubernetes includes centralized or distributed storage shared across nodes in a cluster. For example, this includes Persistent Volumes, which are provisioned by an administrator for storage across Pods and/or containers across all nodes within a Kubernetes cluster.</p> <p><i>A <b>PersistentVolume</b> (PV) is a piece of storage in the cluster that has been provisioned by an administrator or dynamically provisioned using <a href="#">Storage Classes</a>. It is a resource in the cluster just like a node is a cluster resource. PVs are volume plugins like Volumes, but have a lifecycle independent of any individual Pod that uses the PV. This API object captures the details of the implementation of the storage, be that NFS, iSCSI, or a cloud-provider-specific storage system.</i></p> <p>Source: <a href="https://kubernetes.io/docs/concepts/storage/persistent-volumes/">https://kubernetes.io/docs/concepts/storage/persistent-volumes/</a>.</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<h2>StorageClass objects</h2> <p>Each StorageClass contains the fields <code>provisioner</code>, <code>parameters</code>, and <code>reclaimPolicy</code>, which are used when a PersistentVolume belonging to the class needs to be dynamically provisioned to satisfy a PersistentVolumeClaim (PVC).</p> <p>The name of a StorageClass object is significant, and is how users can request a particular class. Administrators set the name and other parameters of a class when first creating StorageClass objects.</p> <p>As an administrator, you can specify a default StorageClass that applies to any PVCs that don't request a specific class. For more details, see the <a href="#">PersistentVolumeClaim concept</a>.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/storage/storage-classes/">https://kubernetes.io/docs/concepts/storage/storage-classes/</a>.</p>
[1.h] wherein the first cluster establishes communications between the set of computing resources of the first cluster and a first gateway communicatively linked between the first cluster and the private communications network;	<p>On information and belief, the American Count IX Systems and Services include a system where the first cluster establishes communications between the set of computing resources of the first cluster and a first gateway communicatively linked between the first cluster and the private communications network.</p> <p>For example, Kubernetes clusters include at least one node, where a node is either a physical or virtual machine comprising a CPU or portion of CPU resources and memory, to run workloads. These clusters are connected to a private network, for example, the American accused network via a gateway, for example, such as Ingress controllers, internal load balancers, and cloud private endpoints. On information and belief, the private network, for example, the American accused network, facilitates container-to-container, Pod-to-Pod, and/or Pod-to-Services communications across multiple clusters for sharing data and handling tasks.</p>


U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p><b>Ingress</b></p> <p>Make your HTTP (or HTTPS) network service available using a protocol-aware configuration mechanism, that understands web concepts like URIs, hostnames, paths, and more. The Ingress concept lets you map traffic to different backends based on rules you define via the Kubernetes API.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/ingress/">https://kubernetes.io/docs/concepts/services-networking/ingress/</a>.</p> <p><b>Terminology</b></p> <p>For clarity, this guide defines the following terms:</p> <ul style="list-style-type: none"><li>• Node: A worker machine in Kubernetes, part of a cluster.</li><li>• <b>Cluster</b>: A set of Nodes that run containerized applications managed by Kubernetes. For this example, and in most common Kubernetes deployments, nodes in the cluster are not part of the public internet.</li><li>• Edge router: A router that enforces the firewall policy for your cluster. This could be a gateway managed by a cloud provider or a physical piece of hardware.</li><li>• <b>Cluster network</b>: A set of links, logical or physical, that facilitate communication within a cluster according to the Kubernetes networking model.</li><li>• Service: A Kubernetes <u>Service</u> that identifies a set of Pods using <u>label</u> selectors. Unless mentioned otherwise, Services are assumed to have virtual IPs only routable within the cluster network.</li></ul> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/ingress/">https://kubernetes.io/docs/concepts/services-networking/ingress/</a>.</p>




U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p><b>What is Ingress?</b></p> <p>Ingress exposes HTTP and HTTPS routes from outside the cluster to <b>services</b> within the cluster. Traffic routing is controlled by rules defined on the Ingress resource.</p> <p>Here is a simple example where an Ingress sends all its traffic to one Service:</p>  <pre> graph LR     client((client)) -.-&gt; Ingress-managed load balancer  Ingress[Ingress]     subgraph cluster         Ingress -- routing rule --&gt; Service[Service]         Service --&gt; Pod1[Pod]         Service --&gt; Pod2[Pod]     end </pre> <p>Figure. Ingress</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/ingress/">https://kubernetes.io/docs/concepts/services-networking/ingress/</a>.</p>
[1.i] wherein the second cluster establishes communications among the set of computing resources of the second cluster and a second gateway communicatively linked between the second cluster and the private communications network;	<p>On information and belief, the American Count IX Systems and Services include a system where the second cluster establishes communications among the set of computing resources of the second cluster and a second gateway communicatively linked between the second cluster and the private communications network.</p> <p>Kubernetes clusters include at least one node, where a node is either a physical or virtual machine comprising a CPU or portion of CPU resources and memory, to run workloads.</p> <p>On information and belief, in Kubernetes, Services provide a way for communication within a second cluster to the private network, for example, the American accused network via a gateway, such as Ingress controllers, internal load balancers, and cloud private endpoints. On information and belief, the private network, for example, the American accused network, facilitates container-to-container, Pod-</p>




U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p>to-Pod, and/or Pod-to-Services communications across multiple clusters for sharing data and handling tasks.</p> <h2>Ingress</h2> <p>Make your HTTP (or HTTPS) network service available using a protocol-aware configuration mechanism, that understands web concepts like URIs, hostnames, paths, and more. The Ingress concept lets you map traffic to different backends based on rules you define via the Kubernetes API.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/ingress/">https://kubernetes.io/docs/concepts/services-networking/ingress/</a>.</p> <h2>Terminology</h2> <p>For clarity, this guide defines the following terms:</p> <ul style="list-style-type: none"> <li>• Node: A worker machine in Kubernetes, part of a cluster.</li> <li>• <b>Cluster</b>: A set of Nodes that run containerized applications managed by Kubernetes. For this example, and in most common Kubernetes deployments, nodes in the cluster are not part of the public internet.</li> <li>• Edge router: A router that enforces the firewall policy for your cluster. This could be a gateway managed by a cloud provider or a physical piece of hardware.</li> <li>• <b>Cluster network</b>: A set of links, logical or physical, that facilitate communication within a cluster according to the Kubernetes networking model.</li> <li>• Service: A Kubernetes <u>Service</u> that identifies a set of Pods using <u>label</u> selectors. Unless mentioned otherwise, Services are assumed to have virtual IPs only routable within the cluster network.</li> </ul> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/ingress/">https://kubernetes.io/docs/concepts/services-networking/ingress/</a>.</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p><b>What is Ingress?</b></p> <p><b>Ingress</b> exposes HTTP and HTTPS routes from outside the cluster to <b>services</b> within the cluster. Traffic routing is controlled by rules defined on the Ingress resource.</p> <p>Here is a simple example where an Ingress sends all its traffic to one Service:</p>  <p>The diagram illustrates the flow of traffic from an external client to a Kubernetes cluster. A grey oval labeled 'client' is connected by a dashed arrow labeled 'Ingress-managed load balancer' to a blue box labeled 'Ingress'. The 'Ingress' box is connected by a solid arrow labeled 'routing rule' to a blue box labeled 'Service'. The 'Service' box is connected by two curved arrows to two separate blue boxes labeled 'Pod'. The entire cluster components (Ingress, Service, and Pods) are enclosed in a red rectangular border labeled 'cluster' at the top.</p> <p>Figure. Ingress</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/ingress/">https://kubernetes.io/docs/concepts/services-networking/ingress/</a>.</p> <p>A Kubernetes cluster network includes the Gateway API, which provides Gateway resources. These resources can be defined and specified so that external traffic from the Internet can be routed to various Services.</p>


U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p><b>Gateway API</b></p> <p>Gateway API is a family of API kinds that provide dynamic infrastructure provisioning and advanced traffic routing.</p> <p>Make network services available by using an extensible, role-oriented, protocol-aware configuration mechanism. Gateway API is an add-on containing API kinds that provide dynamic infrastructure provisioning and advanced traffic routing.</p> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/gateway/">https://kubernetes.io/docs/concepts/services-networking/gateway/</a>.</p> <ul style="list-style-type: none"><li>• The Gateway API (or its predecessor, Ingress) allows you to make Services accessible to clients that are outside the cluster.</li></ul> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/">https://kubernetes.io/docs/concepts/services-networking/</a>.</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<h2>Resource model</h2> <p>Gateway API has three stable API kinds:</p> <ul style="list-style-type: none"> <li>• <b>GatewayClass:</b> Defines a set of gateways with common configuration and managed by a controller that implements the class.</li> <li>• <b>Gateway:</b> Defines an instance of traffic handling infrastructure, such as cloud load balancer.</li> <li>• <b>HTTPRoute:</b> Defines HTTP-specific rules for mapping traffic from a Gateway listener to a representation of backend network endpoints. These endpoints are often represented as a <u>Service</u>.</li> </ul> <p>Gateway API is organized into different API kinds that have interdependent relationships to support the role-oriented nature of organizations. A Gateway object is associated with exactly one GatewayClass; the GatewayClass describes the gateway controller responsible for managing Gateways of this class. One or more route kinds such as HTTPRoute, are then associated to Gateways. A Gateway can filter the routes that may be attached to its listeners, forming a bidirectional trust model with routes.</p> <p>The following figure illustrates the relationships of the three stable Gateway API kinds:</p>  <pre> graph LR     subgraph cluster         HTTPRoute --&gt; Gateway         Gateway --&gt; GatewayClass     end </pre> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/gateway/">https://kubernetes.io/docs/concepts/services-networking/gateway/</a>.</p>


U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	 <pre> graph LR     client([client]) -.-&gt; "HTTP request"  Gateway[Gateway]     Gateway --&gt; HTTPRoute[HTTPRoute]     HTTPRoute -- "Routing rule" --&gt; Service[Service]     Service --&gt; Pod1[Pod]     Service --&gt; Pod2[Pod] </pre> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/gateway/">https://kubernetes.io/docs/concepts/services-networking/gateway/</a>.</p> <h2>Design principles</h2> <p>The following principles shaped the design and architecture of Gateway API:</p> <ul style="list-style-type: none"> <li>• <b>Role-oriented:</b> Gateway API kinds are modeled after organizational roles that are responsible for managing Kubernetes service networking:             <ul style="list-style-type: none"> <li>◦ <b>Infrastructure Provider:</b> Manages infrastructure that allows multiple isolated clusters to serve multiple tenants, e.g. a cloud provider.</li> <li>◦ <b>Cluster Operator:</b> Manages clusters and is typically concerned with policies, network access, application permissions, etc.</li> <li>◦ <b>Application Developer:</b> Manages an application running in a cluster and is typically concerned with application-level configuration and Service composition.</li> </ul> </li> <li>• <b>Portable:</b> Gateway API specifications are defined as custom resources and are supported by many implementations.</li> <li>• <b>Expressive:</b> Gateway API kinds support functionality for common traffic routing use cases such as header-based matching, traffic weighting, and others that were only possible in Ingress by using custom annotations.</li> <li>• <b>Extensible:</b> Gateway allows for custom resources to be linked at various layers of the API. This makes granular customization possible at the appropriate places within the API structure.</li> </ul> <p>Source: <a href="https://kubernetes.io/docs/concepts/services-networking/gateway/">https://kubernetes.io/docs/concepts/services-networking/gateway/</a>.</p>



U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
[1.j] wherein communications between the first cluster and the second cluster are isolated;	<p>On information and belief, the American Count IX Systems and Services include a system where communications between the first cluster and the second cluster are isolated.</p> <p>For example, multi-cluster Kubernetes provides for isolated workloads in separate clusters.</p> <p><b>Workload Isolation</b></p> <p>Multi-cluster Kubernetes allows you to isolate different workloads, such as development, staging, and production environments, in separate clusters. This improves fault tolerance and reduces the risk of cascading failures. Workload isolation also enables you to enforce strict resource quotas and security policies on a per-cluster basis.</p> <p>Source: <a href="https://www.tigera.io/learn/guides/kubernetes-networking/kubernetes-multi-cluster/">https://www.tigera.io/learn/guides/kubernetes-networking/kubernetes-multi-cluster/</a>.</p> <p>Multi-cluster Kubernetes also provides for configuring access to different clusters.</p> <p><b>Configure Access to Multiple Clusters</b></p> <p>This page shows how to configure access to multiple clusters by using configuration files. After your clusters, users, and contexts are defined in one or more configuration files, you can quickly switch between clusters by using the <code>kubect1 config use-context</code> command.</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<h2>Define clusters, users, and contexts</h2> <p>Suppose you have two clusters, one for development work and one for test work. In the <code>development</code> cluster, your frontend developers work in a namespace called <code>frontend</code>, and your storage developers work in a namespace called <code>storage</code>. In your <code>test</code> cluster, developers work in the default namespace, or they create auxiliary namespaces as they see fit. Access to the development cluster requires authentication by certificate. Access to the test cluster requires authentication by username and password.</p> <p>Source: <a href="https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/">https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/</a>.</p>
[1.k] wherein the first cluster is a high performance cluster; and	<p>On information and belief, the American Count IX Systems and Services include a system where the first cluster is a high-performance cluster.</p> <p>For example, Kubernetes is a system for automating deployment, scaling and management or containerized applications.</p> <p>Kubernetes, also known as K8s, is an open source system for automating deployment, scaling, and management of containerized applications.</p> <p>It groups containers that make up an application into logical units for easy management and discovery. Kubernetes builds upon 15 years of experience of running production workloads at Google, combined with best-of-breed ideas and practices from the community.</p> <p>Source: <a href="https://kubernetes.io/">https://kubernetes.io/</a>.</p> 



U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<div></div> <div><p>Planet Scale <a href="#">↗</a></p><p>Designed on the same principles that allow Google to run billions of containers a week, <b>Kubernetes can scale</b> without increasing your operations team.</p></div> <p>Source: <a href="https://kubernetes.io/">https://kubernetes.io/</a>.</p> <p><b>Kubernetes is a scalable and performant engine that orchestrates containers in a server environment. It is highly optimized by default, and it scales nicely in a suitable infrastructure.</b></p> <p>It is also less opinionated by default, and there are plenty of customizations for end-users to define. This flexibility allows Kubernetes to cover many different use cases and penetrate the market faster, making it extremely popular.</p> <p>Source: <a href="https://platform9.com/blog/10-kubernetes-performance-tips/">https://platform9.com/blog/10-kubernetes-performance-tips/</a>.</p> <p>As a non-limiting example, HPC workloads include many tasks, servers, and parallelization of processing deployed via Kubernetes.</p>

U.S. Patent No. 8,352,584 (Claim 1)	
Claim 1	Example American Count IX Systems and Services
	<p>In Kubernetes, the base unit of scheduling is a Pod: one or more Docker containers scheduled to a cluster host. Kubernetes assumes that workloads are containers. While Kubernetes has the notion of <a href="#">Cron Jobs</a> and <a href="#">Jobs</a> that run to completion, applications deployed on Kubernetes are typically long-running services, like web servers, load balancers or data stores and while they are highly dynamic with pods coming and going, they differ greatly from HPC application patterns.</p> <p>Traditional HPC applications often exhibit different characteristics:</p> <ul style="list-style-type: none"> <li>• In financial or engineering simulations, a job may be comprised of tens of thousands of short-running tasks, demanding low-latency and high-throughput scheduling to complete a simulation in an acceptable amount of time.</li> <li>• A computational fluid dynamics (CFD) problem may execute in parallel across many hundred or even thousands of nodes using a message passing library to synchronize state. This requires specialized scheduling and job management features to allocate and launch such jobs and then to checkpoint, suspend/resume or backfill them.</li> <li>• Other HPC workloads may require specialized resources like GPUs or require access to limited software licenses. Organizations may enforce policies around what types of resources can be used by whom to ensure projects are adequately resourced and deadlines are met.</li> </ul> <p>Source: <a href="https://kubernetes.io/blog/2017/08/kubernetes-meets-high-performance/#:~:text=HPC%20workloads%20unique%20challenges">https://kubernetes.io/blog/2017/08/kubernetes-meets-high-performance/#:~:text=HPC%20workloads%20unique%20challenges</a>.</p>

[1.1] wherein the second cluster is a high performance cluster.

On information and belief, the American Count IX Systems and Services include a system where the second cluster is a high performance cluster.

For example, Kubernetes is a system for automating deployment, scaling and management or containerized applications.

Kubernetes, also known as K8s, is an open source system for automating deployment, scaling, and management of containerized applications.

It groups containers that make up an application into logical units for easy management and discovery. Kubernetes builds upon 15 years of experience of running production workloads at Google, combined with best-of-breed ideas and practices from the community.



Source: <https://kubernetes.io/>.



Planet Scale

Designed on the same principles that allow Google to run billions of containers a week, Kubernetes can scale without increasing your operations team.

Source: <https://kubernetes.io/>.

Kubernetes is a scalable and performant engine that orchestrates containers in a server environment. It is highly optimized by default, and it scales nicely in a suitable infrastructure.

It is also less opinionated by default, and there are plenty of customizations for end-users to define. This flexibility allows Kubernetes to cover many different use cases and penetrate the market faster, making it extremely popular.

Source: <https://platform9.com/blog/10-kubernetes-performance-tips/>.

As a non-limiting example, HPC workloads include many tasks, servers, and parallelization of processing deployed via Kubernetes.

In Kubernetes, the base unit of scheduling is a Pod: one or more Docker containers scheduled to a cluster host. Kubernetes assumes that workloads are containers. While Kubernetes has the notion of [Cron Jobs](#) and [Jobs](#) that run to completion, applications deployed on Kubernetes are typically long-running services, like web servers, load balancers or data stores and while they are highly dynamic with pods coming and going, they differ greatly from HPC application patterns.

Traditional HPC applications often exhibit different characteristics:

- In financial or engineering simulations, a job may be comprised of tens of thousands of short-running tasks, demanding low-latency and high-throughput scheduling to complete a simulation in an acceptable amount of time.
- A computational fluid dynamics (CFD) problem may execute in parallel across many hundred or even thousands of nodes using a message passing library to synchronize state. This requires specialized scheduling and job management features to allocate and launch such jobs and then to checkpoint, suspend/resume or backfill them.
- Other HPC workloads may require specialized resources like GPUs or require access to limited software licenses. Organizations may enforce policies around what types of resources can be used by whom to ensure projects are adequately resourced and deadlines are met.

Source: <https://kubernetes.io/blog/2017/08/kubernetes-meets-high-performance/#:~:text=HPC%20workloads%20unique%20challenges>.